A Fast Super-Resolution Reconstruction Algorithm for Pure Translational Motion and Common Space-Invariant Blur

Michael Elad, *Member, IEEE,* and Yacov Hel-Or

***Abstract*—This paper addresses the problem of recovering a super-resolved image from a set of warped blurred and decimated versions thereof. Several algorithms have already been proposed for the solution of this general problem. In this paper, we concentrate on a special case where the warps are pure translations, the blur is space invariant and the same for all the images, and the noise is white. We exploit previous results to develop a new highly efficient super-resolution reconstruction algorithm for this case, which separates the treatment into de-blurring and measurements fusion. The fusion part is shown to be a very simple non-iterative algorithm, preserving the optimality of the entire reconstruction process, in the maximum-likelihood sense. Simulations demonstrate the capabilities of the proposed algorithm.**

***Index Terms*—Maximum-likelihood, reconstruction, super-resolution, translation motion.**

I. Introduction

The super-resolution reconstruction problem is well known and extensively treated in the literature [1]–[13]. The main idea is to recover a single high-resolution image from a set of low quality images of the same photographed object. In this process, it is conceptually possible to remove some of the aliasing and to increase the effective resolution of the sensor up to the optical cut-off frequency. Recent work [10]–[13] relates this problem to restoration theory [14]. As such, the problem is shown to be an inverse problem, where an unknown image is to be reconstructed, based on measurements related to it through linear operators and additive noise. This linear relation is composed of geometric warp, blur and decimation operations. In [13] a solution (using the maximum-likelihood (ML), maximum *a-posteriori* (MAP), and projection onto convex sets (POCS) methods) to the super-resolution reconstruction problem is given in a simple yet general algebraic form. The proposed solution can deal with a general geometric warp, space-varying blur, spatially uniform decimation with rational resolution ratio, and colored Gaussian additive noise. The solution is based on the knowledge of the operators involved and the noise characteristics.

This paper concentrates on a special super-resolution case. It is assumed that the blur is space invariant and the same for all the measured images; the geometric warps between the measured images are pure translations; and the additive noise is white. These assumptions are indeed very limiting, but in some cases are quite practical. Such is the case in video sequences where the photographed scene is static and the images are obtained with slight translations. Another relevant application is increasing a scanner resolution by scanning the same original document several times with slight different initial points. Several papers already dealt with this special case [2], [4], [6], [7] and proposed different reconstruction algorithms.

In this paper, we propose a new algorithm for the above special super-resolution case. The algorithm is based on the general solution proposed in [13]. Exploiting the properties of the operations involved, the well-known fact that the de-blurring can be separated from the fusion process is first established [2], [6]. The main contribution of this paper corresponds to the fusion stage, where the measurements are merged into a higher resolution image. It is shown that through a very simple noniterative algorithm, this fusion is achieved, while preserving the optimality in the Maximum-Likelihood sense. The new algorithm is shown to be superior to the existing algorithms [2]–[13] in terms of computational cost, and with high output quality.

II. General Super-Resolution

In this section, we briefly describe the general super-resolution model and solution. Detailed description of these topics can be found in [13]. We denote the N measured images by $\{\underline{Y}_k\}_{k=1}^N$. These images are to be fused into a single improved quality image, denoted as $\underline{X}$. The images are represented lexicographically ordered column vectors. Each of these images is related to the required super-resolution image through geometric warp, blur, decimation, and additive noise

$$\underline{Y}_k = D_k H_k F_k \underline{X} + \underline{V}_k \quad k = 1, \ldots, N. \tag{1}$$

The matrix F_k stands for the geometric warp operation that exists between the images $\underline{X}$ and an interpolated version of the image $\underline{Y}_k$ (interpolation is required in order to treat the image $\underline{Y}_k$ in the higher resolution grid). The matrix H_k is the blur matrix, representing the camera's PSF. The matrix D_k stands for the decimation operation, representing the reduction of the number of observed pixels in the measured images. The vectors $\{\underline{V}_k\}_{k=1}^N$ represent Gaussian additive measurement noise with zero mean and auto-correlation matrix $W_k = E\{\underline{V}_k \underline{V}_k^T\}$.

Our aim is to estimate $\underline{X}$ based on the known images $\{\underline{Y}_k\}_{k=1}^N$, and the operations they went through. In order to do that we have

Manuscript received July 28, 1999; revised April 9, 2001. This work was performed while the authors were with Hewlett-Packard Laboratories, Israel, in 1998. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Timothy J. Schulz.

M. Elad is with the Jigami Corporation, The Technion City 32000, Israel (e-mail: elad@jigami.com).

Y. Hel-Or is with the Inter-Disciplinary Center, Herzelia, Israel (e-mail: toky@idc.ac.il).

Publisher Item Identifier S 1057-7149(01)06037-7.

to know D_k, H_k, F_k and W_k for all $k = 1, \ldots, N$. F_k is obtained through motion estimation [15] between images $\{\underline{Y}_k\}_{k=1}^N$ and the image $\underline{Y}_1$ (chosen as a reference image). In order to use the geometric displacements in terms of the finer grid, the obtained motion vectors are to be scaled by the resolution ratio. As for the decimation operation, in most applications all the decimation operations are equal—$\forall k, D_k = D$, and D is defined by the number of pixels in the camera detector array, and the resolution ratio we want to obtain. Similarly, it is typically assumed that all the obtained images go through the same PSF, and therefore $\forall k, H_k = H$. In order to determine H, we can either guess it or estimate it [14]. In most applications the noise is assumed to be white, which means that $E\{\underline{V}_k \underline{V}_k^T\} = \sigma^2 I$. Otherwise, the noise characteristics must be estimated as well.

Based on the assumptions that the additive noise vectors $\{\underline{V}_k\}_{k=1}^N$ are Gaussian and mutually independent, the Maximum-Likelihood estimation of $\underline{X}$ is done through the following least-squares expression [13]:

$$\hat{\underline{X}} = \underset{\underline{X}}{\text{ArgMin}} \left\{ \sum_{k=1}^{N} [\underline{Y}_k - D_k H_k F_k \underline{X}]^T W_k^{-1} [\underline{Y}_k - D_k H_k F_k \underline{X}] \right\}. \tag{2}$$

Taking the first derivative of (2) with respect to $\underline{X}$ and equating to zero we get

$$\sum_{k=1}^{N} [D_k H_k F_k]^T W_k^{-1} [\underline{Y}_k - D_k H_k F_k \underline{X}] = 0 \Rightarrow \mathbf{R}\hat{\underline{X}} = \underline{\mathbf{P}}$$

where $\mathbf{R} = \sum_{k=1}^{N} F_k^T H_k^T D_k^T W_k^{-T} D_k H_k F_k$

$$\underline{\mathbf{P}} = \sum_{k=1}^{N} F_k^T H_k^T D_k^T W_k^{-1} \underline{Y}_k. \tag{3}$$

In order to compute $\hat{\underline{X}}$, the linear system $\mathbf{R}\hat{\underline{X}} = \underline{\mathbf{P}}$ must be solved. This system has a unique solution (i.e., nonsingular $\mathbf{R}$) provided that there are sufficient measurements [13], and that the displacements between the images satisfy some independence constraint [1]–[5], [13]. If these requirements are not met, some regularization scheme must be applied [8]–[13].

Assuming that a unique solution exists, solving $\mathbf{R}\hat{\underline{X}} = \underline{\mathbf{P}}$ directly is practically impossible due to its dimensions. If, for example, the size of image $\hat{\underline{X}}$ is 1000×1000 pixels, the matrix $\mathbf{R}$ is a $10^6 \times 10^6$ matrix. Inversion of such a huge matrix can be obtained using iterative algorithms. Note that the actual inverse of $\mathbf{R}$ (namely $\mathbf{R}^{-1}$) is not required but rather the solution of the linear equation $\mathbf{R}\hat{\underline{X}} = \underline{\mathbf{P}}$. Such iterative methods are very common and very efficient [10]–[16]. Among such algorithms, the steepest descent (SD) algorithm is one of the simplest. The SD algorithm suggests the following iterative equation for the solution of $\mathbf{R}\hat{\underline{X}} = \underline{\mathbf{P}}$:

$$\hat{\underline{X}}_{j+1} = \hat{\underline{X}}_j + \mu[\underline{\mathbf{P}} - \mathbf{R}\hat{\underline{\mathbf{X}}}_j] \tag{4}$$

where $\hat{\underline{X}}_0$, the initialization vector, can be any vector. If R is nonsingular, the above algorithm is guaranteed to converge to the unique solution of $\mathbf{R}\hat{\underline{X}} = \underline{\mathbf{P}}$, provided that $\mu > 0$ is small enough [15]. Putting the terms for $\mathbf{R}$ and $\underline{\mathbf{P}}$ from (3) into the above equation we get

$$\hat{\underline{X}}_{j+1} = \hat{\underline{X}}_j + \mu \sum_{k=1}^{N} F_k^T H_k^T D_k^T W_k^{-1} [\underline{Y}_k - D_k H_k F_k \hat{\underline{X}}_j]. \tag{5}$$

III. Super-Resolution—The Special Case

Let us first repeat the special case properties we intend to exploit.

1) All the decimation operations are assumed to be the same, i.e., $\forall k, D_k = D$.
2) All the blur operations are assumed to be the same, i.e., $\forall k, H_k = H$. Moreover, H is assumed to be block circulant, representing a linear space invariant blur [14].
3) All the warp operations correspond to pure translations. Thus, the matrices F_k are all block-circulant as well [14]. Moreover, we assume that F_k is represented through the nearest neighbor displacement paradigm [15], which means that the displacement in the finer grid is rounded and F_k applies only integer translations. This assumption simplifies the analysis and the obtained results. Its implications on the output quality are negligible, since the rounding is done in the finer resolution grid.
4) The additive noise is white and the same for all the measurements, i.e., $\forall_k, W_k = \sigma^2 I$.

Putting these assumptions into (5), we get that the iterative equation becomes[1] :

$$\hat{\underline{X}}_{j+1} = \hat{\underline{X}}_j + \mu \sum_{k=1}^{N} F_k^T H^T D^T [\underline{Y}_k - DHF_k \hat{\underline{X}}_j]. \tag{6}$$

Exploiting the fact that block circulant matrices commute [14], we get that $F_k^T H^T = H^T F_k^T$ and $HF_k = F_k H$. Thus

$$\hat{\underline{X}}_{j+1} = \hat{\underline{X}}_j + \mu H^T \sum_{k=1}^{N} F_k^T D^T [\underline{Y}_k - DF_k H \hat{\underline{X}}_j]. \tag{7}$$

Let us define the blurred super-resolution image by $\hat{\underline{Z}}_j = H\hat{\underline{X}}_j$. Multiplying both sides of Eq. (7) with H we get

$$\begin{aligned}
H\hat{\underline{X}}_{j+1} &= H\hat{\underline{X}}_j + \mu H H^T \sum_{k=1}^{N} F_K^T D^T [\underline{Y}_k - DF_k H \hat{\underline{X}}_j] \\
\Rightarrow \hat{\underline{Z}}_{j+1} &= \hat{\underline{Z}}_j + \mu H H^T \sum_{k=1}^{N} F_K^T D^T [\underline{Y}_k - DF_K \hat{\underline{Z}}_j] \\
&= \hat{\underline{Z}}_j + \mu H H^T \left[\sum_{k=1}^{N} F_k^T D^T \underline{Y}_k - \sum_{k=1}^{N} F_k^T D^T D F_k \hat{\underline{Z}}_j \right] \\
&= \hat{\underline{Z}}_j + \mu H H^T [\tilde{\underline{\mathbf{P}}} - \tilde{\mathbf{R}} \hat{\underline{Z}}_j]
\end{aligned} \tag{8}$$

where we have used the notations

$$\tilde{\underline{\mathbf{P}}} = \sum_{k=1}^{N} F_k^T D^T \underline{Y}_k \quad \text{and} \quad \tilde{\mathbf{R}} = \sum_{k=1}^{N} F_k^T D^T D F_k. \tag{9}$$

[1]The constant σ is absorbed by the parameter μ

Since the matrix HH^T is positive semi-definite, the above iterative equation stands as a general gradient descent algorithm with a weight matrix HH^T [16]. It is known [16] that such an iterative equation converges to the same final solution as the one without the weight matrix, as long as this matrix is positive semi-definite. Therefore, the steady-state solution of the difference (8) is given $\underline{\hat{Z}}_\infty = \tilde{\mathbf{R}}^{-1}\tilde{\underline{\mathbf{P}}}$. In order to be precise, since the matrix H is typically singular, the steady-state solution consists of two parts: the first is the part of the initialization vector $\underline{\hat{Z}}_0$, which is in the null-space of HH^T, and the second is the solution $\tilde{\mathbf{R}}^{-1}\tilde{\underline{\mathbf{P}}}$, which is orthogonal to this null-space.

Assuming that we somehow found $\underline{\hat{Z}}_\infty$, the above analysis implies that an image restoration process must be applied in order to remove the effect of the blur matrix H. This way, based on $\underline{\hat{Z}}_\infty = H\underline{\hat{X}}_\infty$, we recover the required image $\underline{\hat{X}}_\infty$. The fact that the treatment of the blur can be separated from the fusion of the measurements part was already proposed in other work [2], [3], [6]. For nonsingular matrix H we can say that this process, of first finding a blurred version of the super-resolution image and later restoring the image itself, is as optimal as the direct approach. Since H is typically singular, some regularization element must be added into the recovery process, as we shall show in the following.

We return now to the recovery of the image $\underline{\hat{Z}}_\infty$. As it turns out, computing $\underline{\hat{Z}}_\infty$ is very easy because of the following result—Based on the above assumptions, the matrix $\tilde{\mathbf{R}} = \sum_{k=1}^{N} F_k^T D^T D F_k$ turns out to be a diagonal matrix. In the Appendix we prove this property. Since the matrix $\tilde{\mathbf{R}}$ is diagonal, obtaining $\underline{\hat{Z}}_\infty = \tilde{\mathbf{R}}^{-1}\tilde{\underline{\mathbf{P}}}$ is easy to achieve. The super-resolution reconstruction process thus consists of the following stages.

1) Compute the pair $\tilde{\underline{\mathbf{P}}}$ and $\tilde{\mathbf{R}}$ based on (9). Both of them are stored as images of the same size as the image $\underline{X}$. Note that their computation requires only additions as both the entries of D and F_k are zeros and ones.
2) Compute $\underline{\hat{Z}}_\infty = \tilde{\mathbf{R}}^{-1}\tilde{\underline{\mathbf{P}}}$. This operation requires only one division per pixel. One possible problem with this stage is the possibility that some of the entries on the diagonal of $\tilde{\mathbf{R}}$ may be zero. Such situations occur if there insufficient measurements, or if these measurements are dependent [1]–[5]. In these zero positions, the division $\underline{\hat{Z}}_\infty = \tilde{\mathbf{R}}^{-1}\tilde{\underline{\mathbf{P}}}$ becomes singular. As it turns out, however, it is easy to verify that whenever $\tilde{\mathbf{R}}(m,m) = 0$, it is guaranteed that $\tilde{\mathbf{P}}(m) = 0$, and thus, we get that $\underline{\hat{Z}}_\infty(m) = 0/0$. Some interpolation can be used to fill-in these positions. This interpolation plays a roll of regularization in our recovery process, where we regard the equation $\tilde{\mathbf{R}}\underline{\hat{Z}} = \tilde{\underline{\mathbf{P}}}$ as a constraint, and seek for maximally smoothed solution. Better/different regularization paradigms may be applied in order to obtain edge-preserving behavior [14].
3) Restore $\underline{\hat{X}}_\infty$ from $\underline{\hat{Z}}_\infty$, which can be done in various ways [14]. This part of the process is the computationally demanding part.

There is a simple and intuitive interpretation for the values on the diagonal of $\tilde{\mathbf{R}}$ and the vector $\tilde{\underline{\mathbf{P}}}$. It is easy to show that each diagonal entry in $\tilde{\mathbf{R}}$ corresponds to one pixel in the super-resolution image. Its value is a nonnegative integer, counting the number of measurements contributing to it. The vector $\tilde{\underline{\mathbf{P}}}$ is simply the addition of the measurements after proper zero-filling interpolation and motion compensation. Thus, $\underline{\hat{Z}}_\infty = \tilde{\mathbf{R}}^{-1}\tilde{\underline{\mathbf{P}}}$ is none-other than the pixel-wise average of the measurements.

IV. Relation to Other Methods

Several papers addressed the general super-resolution problem and suggested practical reconstruction algorithms for solving it [8]–[13]. These include the IBP method [10], [11], the POCS-based solution [8], [9], [13], and the MAP based algorithms [12], [13]. These algorithms typically tend to be complex, as they attempt to treat the general problem. In [10], [11], it was suggested to initialize the iterative super-resolution reconstruction process by the image

$$\frac{1}{N}\sum_{k=1}^{N} F_k^{-1} S \underline{Y}_k = \frac{1}{N}\sum_{k=1}^{N} F_k^T S \underline{Y}_k$$

where S is a bilinear interpolator. As it turns out, for the special case treated here, this is not far from the ML optimal solution. There are, however, two important differences: 1) the proper interpolation to use is D^T, which is also simpler and 2) The averaging should be done in a pixel-wise manner, as $\tilde{\mathbf{R}}^{-1}$ does.

The algorithm proposed in this paper performs interpolation and fusion of the different images into a single canvas. This idea was already proposed in [7] based on intuition only. The novelty of our work is in the theoretic justification of such an interpolation-fusion scheme. When facing the special simplified super-resolution problem treated in this paper, one can use one of the following three options:

1) frequency domain methods [2], [3];
2) generalized sampling theorems [4]–[6];
3) interpolation based methods, as the one described in this paper and in [7].

In these three approaches, the ability to separate the treatment of the blur from the fusion of the images can be (and actually is) exploited.

The frequency approach [2], [3] suggests applying a 2-D DFT to each of the input images, combining the images in the frequency domain, exploiting aliasing relationships, and then applying an inverse 2-D DFT. As in our case, blur treatment is done at the end of the recovery algorithm. One of the benefits of the frequency domain algorithm is its ability to be recursive, i.e., the ability to add more measurements as they come. Actually, similar behavior can be identified in our algorithm, since both $\tilde{\mathbf{P}}$ and $\tilde{\mathbf{R}}$ are computed as a direct summation of terms, which correspond to different measurements.

As for the computational complexity of the frequency domain algorithm, in the nonrecursive approach, the frequency domain algorithm requires the accumulation of a complex matrix of size $N \times r$ per pixel[2] and the inversion of it. The recursive approach requires an inversion of a $r \times r$ matrix, and more multiplications in order to apply the RLS algorithm. We have to remember that above all these computations, the DFT operations must be taken into account. Thus, the frequency domain approach is far more complicated, compared to our way of fusing the measurements.

[2] r is the resolution ratio and N is the number of given images

(a) (b) (c)

Fig. 1. Results of the synthetic test. (a) Reference image, (b) original image, and (c) reconstruction results.

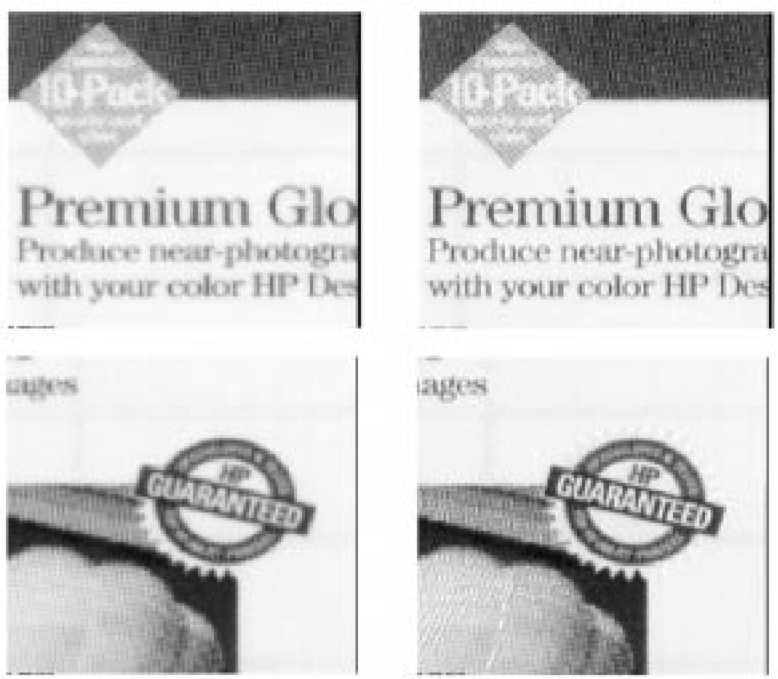

(a) (b)

Fig. 2. Results of the second sequence. (a) Original measured images and (b) result.

The generalized sampling theorems by Yen [4] and later by Papoulis [5] were used as the basis for the method proposed by Ur and Gross [6]. Their method also separates the treatment of the blur from the fusion process. Being based on a sampling theorem, the proposed algorithm is numerically unstable if the measurements have close spatial positions. Similar instability is obtained when there are too many measurements.[3] This algorithm totally disregards the possibility of additive noise. The given samples are considered as the ground truth, and the reconstruction result is merely an interpolation between them.

As for the computational complexity of this algorithm, the recovered signal is computed by summing sufficiently many interpolation functions, which are based on a generalization of the *Sinc* function. In [5] it is claimed that $O\{r^2\}$ multiplications per output pixel are required for the fusion process. Our method, on the other hand, requires only one division per output pixel.

It should be noted that the above complexities refer to the fusion stage only. Assuming that the deblurring part is applied by

[3]Too many implies more than the critical number of required samples.

Fig. 3. Result on the third sequence. (Top) One of the measured images and (bottom) super-resolution result.

Wiener filter, both [5], [7] and our algorithm require a convolution with a relatively large kernel (from 7×7 to 15×15). In the frequency approach, the deconvolution can be done directly on the frequency signal, then requiring an additional one multiplication per one pixel, however, the PSF must be transformed as well.

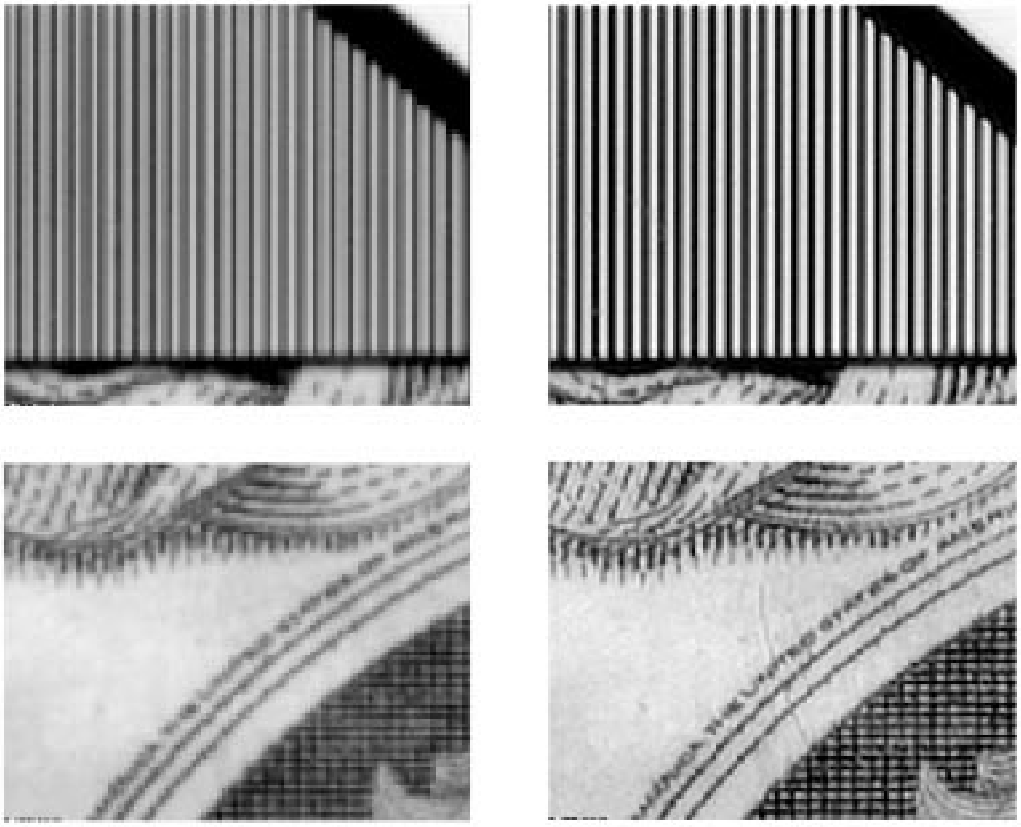

(a) (b)

Fig. 4. Results of the third sequence; sections taken from the full images given in Fig. 3. (a) Measured image and (b) result.

V. RESULTS

We start with a simple synthetic example. We have taken an image of size 720 $\times$ 884 pixels, and created from it nine different 240 $\times$ 294 images by 3 : 1 decimation at each axis and starting at the nine possible different locations. Each of these images is shifted by an integer multiplication of 1/3 at each axis, and these displacements are exactly known. Furthermore, by simply interlacing these images together, we get the original image, which stands for 3 : 1 resolution improvement result. We have applied the reconstruction process on these nine images. The displacements were estimated using an algorithm described in [14], and were found to be (after the rounding for the nearest neighbor) the exact ones. Thus, the vector $\underline{\tilde{\mathbf{P}}}$ consists of exactly the required image, and the main diagonal of $\tilde{\mathbf{R}}$ is actually constant and equals to 1. Thus $\underline{\hat{Z}}_\infty = \tilde{\mathbf{R}}^{-1}\underline{\tilde{\mathbf{P}}}$ is the exact super-resolution image. We have assumed a 2-D separable PSF kernel built from the 1D blur $h = [0.25, 0.5, 0.25]$. A Wiener filter of size 15 $\times$ 15 was applied for restoration, manually searching for the best parameter in this filter.

Fig. 1 presents the results for the synthetic example. We have chosen to show two informative blocks from the input, the output and the original images. The input blocks are taken from the reference image (we could take it from any of the images—the quality is the same), and scaled up by a factor 3 using a NN interpolation. The output and the original blocks correspond to the same portions of the image, and the improvement is self-evident.

Figs. 2–4 correspond to simulations on real sequences acquired by a digital camera. In these two cases, there is no reference image to compare to, as done before. Fig. 2 presents the results for a sequence of 12 images of size 318×411. The resolution is increased by a factor 2. Again, two blocks are shown for comparison, and the improvement is evident—the result is sharper and with more details. The original is a dither image, which explains the textured regions, both in the photographed image and the super-resolution result. Figs. 3 and 4 correspond to the second real test. This sequence contains 16 images of size 300 $\times$ 303 pixels. The resolution is increased by a factor 2 in each axis. Fig. 4 presents magnified small portions of them for better comparison. The PSF in this and the next cases were assumed to be Gaussian blur operations with manually found variance.

VI. CONCLUSION

In this paper, we have presented a new algorithm for super-resolution reconstruction, for the special case were the geometric warp between the given images consist of pure translation, the blur is the same for all the measurements, and is space invariant, and the additive noise is white. The proposed algorithm is shown to be very efficient in terms of computational cost, compared to other algorithms. Simulation results demonstrate its capabilities in terms of output quality.

APPENDIX

In order to show that $\tilde{\mathbf{R}}$ is diagonal, we first note that the matrix $D^T D$ is diagonal. This is easily verified by noticing that the operation $D^T D$ stands for decimation followed by interpolation. Thus, if D decimate by factor of r, applying $D^T D$ causes all the positions $[1+mr, 1+nr]$[4] for integer $[m, n]$ to stay unchanged, whereas the remaining pixels are replaced by zeros.

Let us look at the expression $\tilde{\mathbf{R}}_k = F_k^T D^T D F_k$ for some k. We use the notation $f_k(j)$ to denote the jth column of the matrix F_k. If the displacement vector represented by this matrix is $[d_x, d_y]$, $f_k(j)$ will have "1" value at the position[5] $p\{j, d_x, d_y\}$, and zeros elsewhere. In the general case, the $[m, n]$ entry of the matrix $\tilde{\mathbf{R}}_k$ is given by $f_k^T(m) D^T D f_k(n)$. Since $D^T D$ is diagonal, if $m \neq n$, we get that $p\{m, d_x, d_y\} \neq p\{n, d_x, d_y\}$, and thus, $\tilde{\mathbf{R}}_k[m, n] = 0$. If $m = n$, we get that $\tilde{\mathbf{R}}_k[m, m] = x$. Since $\tilde{\mathbf{R}} = \sum_k \tilde{\mathbf{R}}_k$, we get that $\tilde{\mathbf{R}}$ is diagonal as well, and the claim is proved.

ACKNOWLEDGMENT

The authors would like to thank the anonymous reviewers for their helpful comments, which helped improving this manuscript.

REFERENCES

[1] T. S. Huang and R. Y. Tsay, "Multiple frame image restoration and registration," in *Advances in Computer Vision and Image Processing*, T. S. Huang, Ed. Greenwich, CT: JAI, 1984, vol. 1, pp. 317–339.

[2] S. P. Kim, N. K. Bose, and H. M. Valenzuela, "Recursive reconstruction of high resolution image from noisy undersampled multiframes," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. 38, pp. 1013–1027, June 1990.

[3] N. K. Bose, H. Kim, and B. Zhou, "Performance analysis of the TLS algorithm for image reconstruction from a sequence of undersampled noisy and blurred frames," in *Proc. Int. Conf. Image Processing*, Austin, TX, Nov. 1994, pp. 571–575.

[4] L. J. Yen, "On nonuniform sampling of bandwidth limited signals," *IRE Trans. Circuits Theory*, vol. 3, pp. 251–257, Apr. 1956.

[5] A. Papoulis, "Generalized sampling expansion," *IEEE Trans. Circuits Syst.*, vol. CAS-24, pp. 652–654, Nov. 1977.

[6] H. Ur and D. Gross, "Improved resolution from subpixel shifted pictures," in *CVGIP: Graph., Models, Image Process.*, vol. 54, Mar. 1992, pp. 181–186.

[7] J. C. Gillete, T. M. Stadtmiller, and R. C. Hardie, "Aliasing reduction in staring infrared images utilizing subpixel techniques," *Opt. Eng.*, vol. 34, no. 11, p. 3130, Nov. 1995.

[8] A. M. Tekalp, M. K. Ozkan, and M. I. Sezan, "High resolution image reconstruction from low-resolution image sequences, and space varying image restoration," in *IEEE Int. Conf. Acoustics, Speech, Signal Processing*, vol. 3, San Francisco, CA, Mar. 1992, pp. 169–172.

[4]The image begins at row and column 1.

[5]This mapping is one-to-one, i.e., given the motion vector and p, one can compute the value j.

[9] A. J. Patti, M. I. Sezan, and A. M. Tekalp, "Superresolution video reconstruction with arbitrary sampling lattices and nonzero aperture time," *IEEE Trans. Image Processing*, vol. 6, pp. 1064–1076, Aug. 1997.

[10] M. Irani and S. Peleg, "Improving resolution by image registration," *CVGIP: Graph., Models, Image Process.*, vol. 53, pp. 231–239, May 1991.

[11] ——, "Motion analysis for image enhancement: Resolution, occlusion, and transparency," *J. Vis. Commun. Image Represent.*, vol. 4, pp. 324–335, Dec. 1993.

[12] R. R. Schultz and R. L Stevenson, "Extraction of high-resolution frames from video sequences," *IEEE Trans. Image Processing*, vol. 5, pp. 996–1011, June 1996.

[13] M. Elad and A. Feuer, "Restoration of a single superresolution image from several blurred, noisy and undersampled measured images," *IEEE Trans. Image Processing*, vol. 6, pp. 1646–1658, Dec. 1997.

[14] R. L. Lagendijk and J. Biemond, *Iterative Identification and Restoration of Images*. New York: Kluwer, 1991.

[15] J. L. Barron, D. J. Fleet, and S. Beauchemin, "Performance of optical flow techniques," *Int. J. Comput. Vis.*, vol. 12, pp. 43–77, 1994.

[16] D. Bertsekas, *Nonlinear Programming*. Belmont, MA: Athena, 1995.

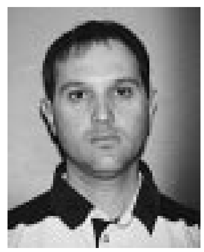

Michael Elad (M'98) was born in Haifa, Israel, in December 1963. He received the B.Sc., M.Sc., and Ph.D. degrees from the Electrical Engineering Department, The Technion—Israel Institute of Technology, Haifa, in 1986, 1988, and 1997, respectively.

From 1997 to 1999, he was with Hewlett-Packard Laboratories, Israel, as a Research and Development Engineer. Since February 2000, he has been with Jigami Corporation, The Technion City, Israel, heading the research division. His current research interests include image reconstruction problems, image and video coding, classification methods for pattern detection and recognition, and optimization theory in the context of image processing, computer vision, and pattern recognition.

Yacov Hel-Or received the B.Sc. degree in physics and computer science from Bar-Ilan University, Israel, in 1985, and the Ph.D. degree in computer science from the Hebrew University, Jerusalem, Israel, in 1993.

During 1993–1994, he was a Postdoctorate Fellow with the Department of Applied Mathematics and Computer Science, Weizmann Institute of Science, Rehovot, Israel. From 1994 to 1996, he was with the NASA Ames Research Center, Moffet Field, CA, as a National Research Council Associate. During these years, he was also affiliated with the Psychology Department, Stanford University, Stanford, CA. From 1996 to 1998, he was with the Hewlett-Packard Science Center, The Technion—Israel Institute of Technology, Haifa. Since 1998, he has been a Faculty Member with the Department of Computer Science, The Inter-Disciplinary Center, Herzelia, Israel. His fields of interests are fields of computer vision, image processing, robotics, and graphics.